Search Results

Search Results for: [pointer and network<AND>((tree <near> search <AND>((direct <near> table<AND>((full <near> match<AND>((longest <near> match))))))))]
Found 24 of 96,434 searched. → Rerun within the Portal

Search within Results

> Advanced Search > Search Help/Tips

Sort by: **Title** **Publication** **Publication Date** **Score** Binder

Results 1 - 20 of 24 **short listing**

Prev Page **1** **2** Next Page

1 Scalable high-speed prefix matching 95%
Marcel Waldvogel , George Varghese , Jon Turner , Bernhard Plattner
ACM Transactions on Computer Systems (TOCS) November 2001
Volume 19 Issue 4

Finding the longest matching prefix from a database of keywords is an old problem with a number of applications, ranging from dictionary searches to advanced memory management to computational geometry. But perhaps today's most frequent best matching prefix lookups occur in the Internet, when forwarding packets from router to router. Internet traffic volume and link speeds are rapidly increasing; at the same time, a growing user population is increasing the size of routing tables against which p ...

2 Small forwarding tables for fast routing lookups 87%
Mikael Degermark , Andrej Brodnik , Svante Carlsson , Stephen Pink
ACM SIGCOMM Computer Communication Review , Proceedings of the ACM SIGCOMM '97 conference on Applications, technologies, architectures, and protocols for computer communication October 1997
Volume 27 Issue 4

3 High-speed policy-based packet forwarding using efficient multi-dimensional range matching 56%
T. V. Lakshman , D. Stiliadis
ACM SIGCOMM Computer Communication Review , Proceedings of the ACM SIGCOMM '98 conference on Applications, technologies, architectures, and protocols for computer communication October 1998
Volume 28 Issue 4

4 Routing on longest-matching prefixes 37%
Willibald Doeringer , Günter Karjoth , Mehdi Nassehi
IEEE/ACM Transactions on Networking (TON) February 1996
Volume 4 Issue 1

5 External memory algorithms and data structures 36%
Jeffrey Scott Vitter
ACM Computing Surveys (CSUR) June 2001
Volume 33 Issue 2
Data sets in large applications are often too massive to fit completely inside the computers internal memory. The resulting input/output communication (or I/O) between fast internal memory and slower external memory (such as disks) can be a major performance bottleneck. In this article we survey the state of the art in the design and analysis of external memory (or EM) algorithms and data structures, where the goal is to exploit locality in order to reduce the I/O costs. We consider a varie ...

6 Technique for automatically correcting words in text 30%
Karen Kukich
ACM Computing Surveys (CSUR) December 1992
Volume 24 Issue 4
Research aimed at correcting words in text has focused on three progressively more difficult problems:(1) nonword error detection; (2) isolated-word error correction; and (3) context-dependent work correction. In response to the first problem, efficient pattern-matching and n-gram analysis techniques have been developed for detecting strings that do not appear in a given word list. In response to the second problem, a variety of general and application-specific spelling cor ...

7 Graph rewrite systems for program optimization 22%
Uwe Assmann
ACM Transactions on Programming Languages and Systems (TOPLAS) July 2000
Volume 22 Issue 4

Graph rewrite systems can be used to specify and generate program optimizations. For termination of the systems several rule-based criteria are developed, defining exhaustive graph rewrite systems. For nondeterministic systems stratification is introduced which automatically selects single normal forms. To illustrate how far the methodology reaches, parts of the lazy code motion optimization are specified. The resulting graph rewrite system classes can be e ...

8 Router plugins 22%

Dan Decasper , Zubin Dittia , Guru Parulkar , Bernhard Plattner
IEEE/ACM Transactions on Networking (TON) February 2000
Volume 8 Issue 1

9 A guided tour to approximate string matching 20%

Gonzalo Navarro
ACM Computing Surveys (CSUR) March 2001
Volume 33 Issue 1

We survey the current techniques to cope with the problem of string matching that allows errors. This is becoming a more and more relevant issue for many fast growing areas such as information retrieval and computational biology. We focus on online searching and mostly on edit distance, explaining the problem and its relevance, its statistical behavior, its history and current developments, and the central ideas of the algorithms and their complexities. We present a number of experiments to ...

10 Routing with a clue 19%

Yehuda Afek , Anat Bremler-Barr , Sariel Har-Peled
IEEE/ACM Transactions on Networking (TON) December 2001
Volume 9 Issue 6

We suggest a new simple forwarding technique to speed up IP destination address lookup. The technique is a natural extension of IP, requires 5 bits in the IP header (IPv4, 7 in IPv6), and performs IP lookup nearly as fast as IP/Tag switching but with a smaller memory requirement and a much simpler protocol. The basic idea is that each router adds a "clue" to each packet, telling its downstream router where it ended the IP lookup. Since the forwarding tables of neighboring routers are similar, th ...

11 Routing with a clue 18%

Anat Bremler-Barr , Yehuda Afek , Sariel Har-Peled
ACM SIGCOMM Computer Communication Review , Proceedings of the conference on Applications, technologies, architectures, and protocols for computer communication August 1999

Volume 29 Issue 4

12 Copy detection for intellectual property protection of VLSI designs 17%

Andrew B. Kahng , Darko Kirovski , Stefanus Mantik , Miodrag Potkonjak , Jennifer L. Wong

Proceeding of the 1999 international conference on Computer-aided design November 1999

We give the first study of copy detection techniques for VLSI CAD applications; these techniques are complementary to previous watermarking-based IP protection methods in finding and proving improper use of design IP. After reviewing related literature (notably in the text processing domain), we propose a generic methodology for copy detection based on determining basic elements within structural representations of solutions (IPs), calculating (context-independent) signatur ...

13 Data compression with finite windows 11%

E. R. Fiala , D. H. Greene

Communications of the ACM April 1989

Volume 32 Issue 4

Several methods are presented for adaptive, invertible data compression in the style of Lempel's and Ziv's first textual substitution proposal. For the first two methods, the article describes modifications of McCreight's suffix tree data structure that support cyclic maintenance of a window on the most recent source characters. A percolating update is used to keep node positions within the window, and the updating process is shown to have constant amortized cost. Other methods explore the ...

14 Modeling for text compression 11%

Timothy Bell , Ian H. Witten , John G. Cleary

ACM Computing Surveys (CSUR) December 1989

Volume 21 Issue 4

The best schemes for text compression use large models to help them predict which characters will come next. The actual next characters are coded with respect to the prediction, resulting in compression of information. Models are best formed adaptively, based on the text seen so far. This paper surveys successful strategies for adaptive modeling that are suitable for use in practical text compression systems. The strategies fall into three main classes: finite-context modeling, i ...

15 Scalable high speed IP routing lookups 7%

Marcel Waldvogel , George Varghese , Jon Turner , Bernhard Plattner

ACM SIGCOMM Computer Communication Review , Proceedings of the ACM SIGCOMM '97 conference on Applications, technologies, architectures, and protocols for computer communication October 1997
Volume 27 Issue 4

16 Router plugins: a software architecture for next generation routers 6%
Dan Decasper , Zubin Dittia , Guru Parulkar , Bernhard Plattner
ACM SIGCOMM Computer Communication Review , Proceedings of the ACM SIGCOMM '98 conference on Applications, technologies, architectures, and protocols for computer communication October 1998
Volume 28 Issue 4

17 The J-machine multicomputer 5%
Michael D. Noakes , Deborah A. Wallach , William J. Dally
ACM SIGARCH Computer Architecture News , Proceedings of the 20th annual international symposium on Computer architecture May 1993
Volume 21 Issue 2
The MIT J-Machine multicomputer has been constructed to study the role of a set of primitive mechanisms in providing efficient support for parallel computing. Each J-Machine node consists of an integrated multicomputer component, the Message-Driven Processor (MDP), and 1 MByte of DRAM. The MDP provides mechanisms to support efficient communication, synchronization, and naming. A 512 node J-Machine is operational and is due to be expanded to 1024 nodes in March 1993. In this paper we discuss ...

18 Field programmable port extender (FPX) for distributed routing and queuing 3%
John W. Lockwood , Jon S. Turner , David E. Taylor
Proceedings of the ACM/SIGDA international symposium on Field programmable gate arrays February 2000
Field Programmable Gate Arrays (FPGAs) are being used to provide fast Internet Protocol (IP) packet routing and advanced queuing in a highly scalable network switch. A new module, called the Field-programmable Port Extender (FPX), is being built to augment the Washington University Gigabit Switch (WUGS) with reprogrammable logic.FPX modules reside at the edge of the WUGS switching fabric. Physically, the module is inserted between an optical line card and the WUGS gigabit switch ...

19 The Hearsay-II Speech-Understanding System: Integrating 2%

Knowledge to [redacted]solve Uncertainty
Lee D. Erman , Frederick Hayes-Roth , Victor R. Lesser , D. Raj Reddy
ACM Computing Surveys (CSUR) June 1980
Volume 12 Issue 2

20 Median split trees 2%
B. A. Sheil
Communications of the ACM November 1978
Volume 21 Issue 11
Split trees are a new technique for searching sets of keys with highly skewed frequency distributions. A split tree is a binary search tree each node of which contains two key values—a node value which is a maximally frequent key in that subtree, and a split value which partitions the remaining keys (with respect to their lexical ordering) between the left and right subtrees. A median split tree (MST) uses the lexical median of a node ...

Results 1 - 20 of 24 **short listing**

Prev Page **1** **2** Next Page
